arXiv:1606.05473v1 [cs.DC] 17 Jun 2016

Parallel Reachability Analysis for Hybrid Systems

Amit Gurung[1], Arup Kumar Deka[1],
Ezio Bartocci[2], Sergiy Bogomolov[3], Radu Grosu[2], and Rajarshi Ray[1⋆]

[1] National Institute of Technology Meghalaya, India
[2] Vienna University of Technology, Austria
[3] Institute of Science and Technology Austria, Austria

Abstract. We propose two parallel state-space exploration algorithms for hybrid systems with the goal of enhancing performance on multi-core shared memory systems. The first is an adaption of the parallel breadth first search in the SPIN model checker. We show that the adapted algorithm does not provide the desired load balancing for many hybrid systems benchmarks. The second is a task parallel algorithm based on cheaply precomputing cost of post (continuous and discrete) operations for effective load balancing. We illustrate the task parallel algorithm and the cost precomputation of post operators on a support-function-based algorithm for state-space exploration. The performance comparison of the two algorithms displays a better CPU utilization/load-balancing of the second over the first, except for certain cases. The algorithms are implemented in the model checker XSpeed and our experiments show a maximum speed-up of 900× on a navigation benchmark with respect to SpaceEx LGG scenario, comparing on the basis of equal number of post operations evaluated.

1 Introduction

Hybrid systems are a popular formal framework to model and verify safety properties in biological [2,3] and cyber-physical systems [19]. This formalism combines the classical discrete state-based representation of finite automata with a continuous dynamics semantics in each state (or mode). A hybrid system is called *safe* if a given set of *bad states* are not reachable from a set of initial states. Hence, safety can be proved by performing a reachability analysis that is in general undecidable [15] for hybrid systems. SpaceEx [11,6,5,4] is one of the most popular reachability-analysis tools for hybrid systems using semi-decision procedures based on over-approximation techniques [13,12]. The reachable states are

[0] **Acknowledgements.** This work was supported in part by DST-SERB, GoI under Project No. YSS/2014/000623 and by the European Research Council (ERC) under grant 267989 (QUAREM) and by the Austrian Science Fund (FWF) under grants S11402-N23, S11405-N23 and S11412-N23 (RiSE/SHiNE) and Z211-N23 (Wittgenstein Award).

⋆ Corresponding author e-mail address: `rajarshi.ray@nitm.ac.in`

represented as a collection of continuous sets, each one symbolically represented. The main two challenges to address with such set-based methods are precision and scalability. Recently, algorithms using convex sets represented using support functions [13,14] and zonotopes [12] have shown a very good scalability. However, all these techniques were originally conceived to run sequentially. Hence, they are currently not suitable to exploit the modern and powerful multi-core architectures that would enable them to improve further their scalability.

Our contribution In this work, we provide two parallel algorithms for reachability analysis of hybrid systems that can leverage CPU multi-core architectures to speed-up the performance of the current technology. Our approach relies on the adaptation of Holzmann's lock-free parallel breadth first exploration algorithm [16] recently implemented in the SPIN model checker. We first extend the original algorithm to deal with the symbolic reachable states and flowpipe computations that are the necessary ingredients of reachability analysis of hybrid systems. However, we notice that this first approach often results not ideal concerning the CPU's cores utilization and the load balancing. This happens when the number of symbolic states to be explored is less than the number of processor's cores or when the cost of flowpipe computation varies drastically for different reachable states. For this reason, we provide a second algorithm that improves considerably the load balancing by efficiently precomputing the cost of certain operations. The two algorithms are implemented in XSpeed [22] tool that is now able to handle also SpaceEx models using the Hyst [1] model transformation and translation tool for hybrid automaton models. Our experiments show a speed-up of up to 900× on a navigation benchmark instance with respect to SpaceEx LGG scenario, comparing on the basis of equal number of post operations evaluated. The tool and the benchmarks reported in the paper can be downloaded from `http://nitmeghalaya.in/nitm_web/fp/cse_dept/XSpeed/index_new.html`

Related Work In the last decade, there has been an increasing interest in developing techniques for reachability analysis for hybrid systems. The tools currently [11,7,17,20] available can perform reachability analysis of hybrid systems with linear dynamics and more than 200 variables [11,14]. However, all these works are not suitable to exploit the powerful modern parallel multi-core architectures. In our previous paper [21,22] we addressed the problem of computing reachability analysis of continuous systems, but we did not handle hybrid systems. Hence, to the best of our knowledge, this paper represents the first attempt to provide a parallel reachability analysis algorithm for hybrid systems.

Paper organization The rest of the paper is organized as follows. Section 2 provides the necessary background on hybrid automata and reachability analysis. In Section 3, we show how to extend the Holzmann's lock-free parallel breadth first exploration algorithm [16] to handle the state space exploration in hybrid automata. Section 4 addresses the load balancing problem introducing the notion of a task parallel algorithm. In Section 5 we provide a concrete example of a task parallel algorithm in the context of support-function-based reachability analysis. Section 6 reports the experimental results to illustrate performance speed-up and CPU's core utilization. We conclude in Section 7.

2 Preliminaries

A hybrid automaton is a mathematical model of systems exhibiting both continuous and discrete dynamics. A *state* of a hybrid automaton is an n-tuple $(x_1, x_2, \ldots, x_n)$ representing the values of the n continuous variables in an n dimensional system at an instance of time.

Definition 1. *A hybrid automaton is a tuple* $(\mathcal{L}, \mathcal{X}, Inv, Flow, Init, \delta, \mathcal{G}, \mathcal{M})$ *where:*

- $\mathcal{L}$ *is the set of locations of the hybrid automata.*
- $\mathcal{X}$ *is the set of continuous variables of the hybrid automata.*
- $Inv : \mathcal{L} \to \mathbb{R}^n$ *maps every location of the automata to a subset of* $\mathbb{R}^n$*, called the invariant of the location. An invariant of a location defines a bound on the states within the location of the automata.*
- $Flow$ *is a mapping of the locations of the automata to ODE equations of the form* $\dot{x} = f(x), x \in \mathcal{X}$*, called the flow equation of the location. A flow equation defines the evolution of the system variables within a location.*
- $Init$ *is a tuple* $(\ell_{init}, \mathcal{C}_{init})$ *such that* $\ell_{init} \in \mathcal{L}$ *and* $\mathcal{C}_{init} \subseteq Inv(\ell_{init})$*. It defines the set of initial states of the automata.*
- $\mathcal{G} \subseteq \mathbb{R}^n$ *is the set of guard sets of the automata.*
- $\mathcal{M} : \mathbb{R}^n \to \mathbb{R}^n$ *is the set of assignment maps of the automata.*
- $\delta \subseteq \mathcal{L} \times \mathcal{G} \times \mathcal{M} \times \mathcal{L}$ *is the set of transitions of the automata. A transition from a source location to a destination location in* $\mathcal{L}$ *may trigger when the* state s *of the hybrid automata lies in the guard set from* $\mathcal{G}$*. The map* $\mathcal{M}$ *of the transition transforms the state* s *in the source location to a new state* s' *in the destination location.*

A reachable state is a state attainable at any time instant $0 \leq t \leq T$ under its flow and transition dynamics starting from an initial state in $Init$. The flow dynamics evolves a state (ℓ, x) to another state (ℓ, y) in a location ℓ after T time units such that $flow(x, T) = y$ and $flow(x, t) \in \ Inv(\ell)$ for all $t \in [0, T]$, where $flow$ is the solution to the flow equation of $Flow(\ell)$. Reachability analysis tools produce a conservative approximation of the reachable states of the automaton over a time horizon. Reachable states can be expressed as a union of *symbolic states*. A symbolic state is a tuple $\{loc, \mathcal{C}\}$ such that $loc \in \mathcal{L}$ and $\mathcal{C} \subseteq Inv(loc)$.

In Algorithm 1 [11], we show a generic reachability algorithm for hybrid automata. The algorithm maintains two data structures, $Wlist$ and R. $Wlist$ stores the list of symbolic states to initiate the exploration of reachable states and R stores the already visited reachable states. $Init$ is a symbolic state depicting the initial states given as an input. $Wlist$ and R are initialized to $Init$ and $\emptyset$ respectively in line 2. A symbolic state S is removed from the $Wlist$ at each iteration and explored for reachable states. The operators $PostC$ and $PostD$ are applied consecutively. $PostC$ takes a symbolic state as argument and computes the reachable states under the continuous dynamics of the location. The result of $PostC$ is a symbolic state, say R'. When R' is contained in R, there is no newly explored state and the next symbolic state in $Wlist$ is explored as shown

in line 5. On the contrary, when new states are found in R' not in R, they are included in R in line 6 . The operator $PostD$ takes a symbolic state and returns a set of symbolic states obtained under the discrete dynamics. This new set of symbolic states, shown as R'', is added to the $Wlist$ for further exploration.

Algorithm 1 Reachability Algorithm for Hybrid Automata

```
1: procedure REACH-HA(ha,Init)
2:     Wlist ← Init; R ← ∅;
3:     while Wlist ≠ ∅ do
4:         S ← Wlist.pop(); R' ← PostC(S)
5:         if R' ⊆ R then go to step 3
6:         else  R ← R ∪ R'
7:         end if
8:         R'' ← PostD(R'); Wlist.add(R'')
9:     end while
10: end procedure
```

3 Parallel Breadth First Search

It is worthy noting that the $PostC$ computation for symbolic states in $Wlist$ is independent of each other and therefore can be potentially parallelized. In this work, we exploit this inherent parallelism and propose parallel breadth first search (BFS) algorithm. In a multi-threaded implementation, threads can compute $PostC$ and $PostD$ operations in parallel, however the access to the shared data structure $Wlist$ and R has to be mutually exclusive to avoid race condition. Mutual exclusion can be accomplished by using locks or semaphores, however at the price of additional overhead. Moreover, such an implementation may not ensure an effective load balancing as illustrated later in the following.

In Algorithm 2 we show how to avoid the overhead of the mutual exclusion discipline by adapting the parallel lock-free breadth first search algorithm proposed in [16] to hybrid system state-space exploration. Our adapted algorithm is referred as A-GJH (Adapted Gerard J. Holzmann's) algorithm in the paper. The algorithm uses two copies of $Wlist$, each being a two-dimensional list of symbolic states. At each iteration, symbolic states are read from the $Wlist[t]$ copy and new symbolic states are written to the $Wlist[1-t]$ copy. At the first iteration the value t is 0 and at the end of each iteration (see line 20) of the main while loop it is assigned to $1-t$. In this way, in the next iteration the write list becomes the read list and vice-versa. There are N threads, one thread per core, which computes the post operations in parallel. All the symbolic states present in the row $Wlist[t][w]$ are sequentially processed by the thread indexed by w, shown in line 7-13. A symbolic state when undergoes the $PostD$ operation generates a list of successor symbolic states, to be processed in the next iteration. Each successor state is added to the list $Wlist[1-t][w'][w]$, where

w' is randomly selected between one to N (line 13). This random distribution of new states across rows of $Wlist[1-t]$ is used for load balancing. Since each thread indexed at w writes to the list at $Wlist[1-t][w'][w]$, there is no write contention in $Wlist[1-t]$. Checking containment of symbolic states in the result data structure R is avoided in the algorithm since it is an expensive operation. The exploration is instead bounded by the number of levels of exploration of the automata. The symbolic states of $Wlist[t]$ are explored when all the N threads terminate and synchronize at line 16 ensuring a breadth-first exploration. The algorithm terminates when there are no successor states in $Wlist[1-t]$ for further exploration or when the breadth exploration reaches a certain bound.

Algorithm 2 Adapted G.J. Holzmann's Algorithm

```
 1: procedure REACH-PBFS(ha, Init, bound)
 2:     t = 0, level = 0, N = Cores
 3:     Wlist[2][N][N]                ▷ Wlist is a read/write list of symbolic states.
 4:     Wlist[t][1][1] = Init
 5:     Create w = 1 to N threads                                  ▷ N worker threads
 6:     repeat
 7:         for q = 1 to N do
 8:             for each s in q do
 9:                 delete s from Wlist[t][w][q]
10:                 R'[i] ← PostC(s)
11:                 R''[i] ← PostD(R'[i])
12:                 w' = choose random 1...N
13:                 add s' ∈ R''[i] to Wlist[1 − t][w'][w]
14:             end for
15:         end for
16:         Barrier synchronization          ▷ Thread synchronization to ensure BFS
17:         if Wlist[1 − t][1...N][1...N] is all empty then
18:             done = true;
19:         else
20:             t = 1 − t                                          ▷ Read/Write switching
21:             level ← level + 1
22:         end if
23:     until !done OR level > bound
24: end procedure
```

3.1 Load Balancing

The clever use of the data structures in A-GJH algorithm provides freedom from locks and reasonable load balancing when there are sufficiently large number of symbolic states in the waiting list. However, the load balancing in A-GJH does not perform well when the number of symbolic states in the waiting list is less than the number of cores of the processor. Since the symbolic states are

Models	Breadths	Time in Secs			CPU Utilization (%)		Speed-up	
		SpaceEx (LGG)	XSpeed (Seq-BFS)	XSpeed (A-GJH)	XSpeed (Seq-BFS)	XSpeed (A-GJH)	A-GJH vs. Seq-BFS	A-GJH vs. SpaceEx
Circle	4	37.68	29.39	33.92	12.50	12.53	0.9	1.1
	6	86.26	80.47	62.80	12.50	19.80	1.3	1.4
Nav 3x3 (inst# 1)	3	61.42	9.17	5.72	12.50	23.51	1.6	10.7
	5	199.27	44.37	24.24	12.50	39.28	1.8	8.2
Nav 5x5	5	215.46	50.77	31.39	12.50	33.87	1.6	6.9
	7	1232.31	253.63	104.30	12.50	65.38	2.4	11.8

Fig. 1. Moderate CPU utilization and performance speed-up with A-GJH algorithm.

distributed randomly to the N cores in line 13 for exploration, some of the cores remain idle when there are not enough states to be explored. For this reason there are benchmarks for hybrid systems reachability analysis where an incorrect load balancing results in a low time utilization of the available cores. For example, Figure 1 shows that while the A-GJH running in a 4 core processor provides some improvements in performance compared to SpaceEX LGG (Le Guernic-Girard) and the sequential BFS, the CPU core utilization remain very modest. The best utilization is 65% in the NAV 5X5 benchmark for 7 levels exploration and the worst is 12% in the Circle model. In the Circle model, there are only 2 symbolic states for exploration at each BFS iteration, keeping most of the CPU cores idle.

Another principle problem in load balancing is that the cost of flowpipe computation may vary drastically for different symbolic states. This is illustrated on a 3×3 Navigation benchmark in Figure 2 [9]. The benchmark models the motion of an object in a 2D plane partitioned as a 3×3 grid. Each cell in the grid has a width and height of 1 unit and have a desired velocity v_d. In Figure 2 the cells are numbered from 1 to 9 and the respective desired velocities are shown with a directed vector. Note that there is no desired velocity shown in cell 3 and 7 to distinguish them from the others. Cell 3 is the target whereas cell 7 is the unsafe region. The actual velocity of the object in a cell is given by the differential equation $\dot{v} = A(v - v_d)$, where A is a 2×2 matrix. There is an instantaneous change of dynamics on crossing over to an adjacent cell. The green box is an initial set where the object can start with an initial velocity. The red region shows the reachable states under the hybrid dynamics after a finite number of cell transitions.

Figure 2 shows the reachable states after two and three levels of exploration in Algorithm 2. There are four symbolic states, S_1, S_2, S_3 and S_4, waiting to be explored after two levels of bfs, shown in blue. The symbolic states S_1, S_2 are $\{1, B_1\}$, $\{1, B_2\}$ and S_3, S_4 are $\{5, B_3\}$ and $\{5, B_4\}$ respectively, where 1 and 5 are the location ids and B_1, B_2, B_3 and B_4 are the blue regions in the boundary of location with ids 1 and 4, 1 and 2, 4 and 5, 2 and 5 respectively. Algorithm 2 spawns four threads, one each to compute the flowpipe from the symbolic states.

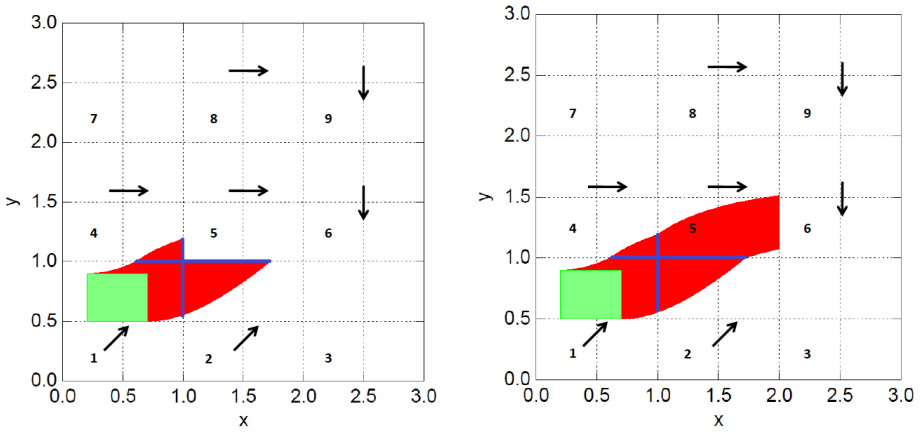

(a) Flowpipe after exploring 2 levels having 4 new symbolic states

(b) Flowpipe after exploring 3 levels

Fig. 2. Illustrating Load Balancing Problem with Flowpipes of Varying Cost

In a four core processor, this seems an ideal load division. However, observe in Figure 2b that out of the four flowpipes, the two from S_1 and S_2 do not lead to new states since they start from the boundary of location id 1 and the dynamics pushes the reachable states outside the location invariant. This implies that the flowpipe computation cost for S_1 and S_2 are low and the two cores assigned to these flowpipe computation finish early and waits at the synchronization point until the remaining two busy cores complete. Such a situation keeps the available cores under-utilized due to the idle waiting of 50% of cores. Similarly, the cost of $PostD$ operation also varies causing idle waiting of threads assigned to low cost computations.

4 Task Parallel Algorithm

In the following we propose an alternative approach to improve load balancing based on computational cost of the tasks encountered during the exploration. The idea is to identify the *atomic tasks* in a flowpipe ($PostC$) computation. The atomic tasks across all flowpipe computations in the breadth search is a measure of the total work-load, at the present breadth of the exploration. For an effective load balancing, this work-load is distributed evenly between the cores of the processor. Similar tasks distribution can be applied also to the computation of discrete transitions ($PostD$). Algorithm 3 shows this approach.

In particular, the instruction in line 7 obtains an estimated cost of computing a flowpipe from a symbolic state using the function $flow_cost$. After the flowpipe costs for all symbolic states in $Wlist$ are computed, line 8 breaks the flowpipe computations into atomic tasks and adds them into a tasks list. In line 12 the

Algorithm 3 Task Parallel Breadth First Exploration

```
1: procedure REACH-TASK-PBFS(ha, Init, bound)
2:     t = 0, level = 0, N = 1, CostC = 0, CostD = 0
3:     Wlist[2][N]                          ▷ Wlist is a read/write list of symbolic states.
4:     Wlist[t][0] = Init
5:     repeat
6:         for each s in Wlist[t] do
7:             CostC = CostC + flow_cost(s)
8:             Break PostC(s) into atomic tasks and add to Tasks list
9:         end for
10:        Tasks_Per_Core = ⌈CostC/#Cores⌉ ▷ Even distribution of tasks to cores
11:        for Threads with id w = 1 to N do                    ▷ N worker threads
12:            Execute Tasks_Per_Core exclusive tasks from the Tasks list
13:            Add results to Res
14:        end for
15:        Barrier Synchronization
16:        for each s in Wlist[t], create thread indexed w do    ▷ w worker threads
17:            flow[w] = Res.join()              ▷ Combine task results to get flowpipe
18:        end for
19:        Barrier Synchronization
20:        for each s in Wlist[t] do
21:            CostD = CostD + jump_cost(flow[s])
22:            Break PostD(s) into atomic tasks and add to Tasks list
23:        end for
24:        Tasks_Per_Core = ⌈CostD/#Cores⌉ ▷ Even distribution of tasks to cores
25:        for Threads with id w = 1 to N do                    ▷ N worker threads
26:            Execute Tasks_Per_Core exclusive tasks from the Tasks list
27:            Add results to Res
28:        end for
29:        for each s in Wlist[t], create thread indexed w do    ▷ w worker threads
30:            R'[w] = Res.join()   ▷ Combine task results to get succ symbolic state
31:            add each s in R'[w] to Wlist[1 − t][w]
32:        end for
33:        Barrier Synchronization
34:        if Wlist[1 − t] is empty then  done = true
35:        else
36:            t = 1 − t                                  ▷ Read/Write switching
37:            N = sum of size of all lists in Wlist[t]
38:            Resize Wlist[1-t][N], level = level + 1, CostC = 0, CostD = 0
39:        end if
40:    until !done OR level > bound
41: end procedure
```

atomic tasks are evenly assigned to the processors cores. In line 17 the results of the atomic tasks computed in parallel are then joined together to obtain a flowpipe. Similar load division is carried out for $PostD$ operation. The successor states obtained from each flowpipe are added in the write list $Wlist[1-t][w]$,

by a thread indexed at w in line 31. The exclusive access of the threads to the columns of $Wlist[1-t]$ eliminate the write contention.

A challenge in this approach is to devise an efficent method for computing the cost of flowpipe and discrete-jump computation for balanced load distribution. Efficient methods and data structures for splitting $PostC$, $PostD$ into atomic tasks and merging their results are very important in order to avoid that the extra required overhead would affect the gain obtained with the parallel exploration. In the next section, we propose some procedures to compute cost of post operations for a support-function algorithm.

5 Task Parallellism in Support Function Algorithm

A common technique in flowpipe computation consists in discretizing the time horizon T into N intervals with a chosen time-step $\tau = T/N$ and over-approximating the reachable states in each interval by a closed convex set, say Ω. A flowpipe can be represented as a union of convex sets as shown in Equation 1.

$$\begin{aligned} &Reach_{[0,T]}(\mathcal{X}_0) \subseteq \bigcup_{i=0}^{N-1} \Omega_i \\ &Reach_{[i\tau,(i+1)\tau]}(\mathcal{X}_0) \subseteq \Omega_i,\ \forall 0 \leq i < N \end{aligned} \tag{1}$$

The representation of the convex sets Ω is a key in the efficiency and scalability of reachability algorithms. Polytopes [8,10] and zonotopes[12] are popular geometric objects suitable to represent convex sets. More recently, support functions [13,14] have been proposed as a functional representation of compact convex sets. SpaceEx [11] and XSpeed [22] tools implement support-function-based algorithms for flowpipe computation due to the promise it provides in scalability. We now present the preliminaries of support functions necessary to introduce the task parallelism in the algorithm.

5.1 Support functions

Definition 2. *[23] Given a nonempty compact convex set $\mathcal{X} \subset \mathbb{R}^n$ the* support function *of $\mathcal{X}$ is a function $sup_{\mathcal{X}} : \mathbb{R}^n \to \mathbb{R}$ defined as:*

$$sup_{\mathcal{X}}(\ell) = \sup\{\ell \cdot x \mid x \in \mathcal{X}\} \tag{2}$$

Definition 3. *Given a support function $sup_{\mathcal{X}}$ of a compact convex set $\mathcal{X}$ and a finite set of template directions $\mathcal{D}$, A template polytope of the convex set $\mathcal{X}$ is defined as:*

$$Poly_{\mathcal{D}}(\mathcal{X}) = \bigcap_{l_i \in \mathcal{D}} l_i.x \leq sup_{\mathcal{X}}(\ell_i) \tag{3}$$

The support-function algorithm in [13] proposes a flowpipe computation by computing the template polyhedral approximation of the convex sets Ω by sampling their support functions in the template directions. The algorithm is for linear dynamics with non-deterministic inputs of the form:

$$\dot{x} = Ax(t) + u(t), \qquad u(t) \in \mathcal{U}, x(0) \in \mathcal{X}_0 \tag{4}$$

where $x(t), u(t) \in \mathbb{R}^n$, A is a real-valued $n \times n$ matrix and $\mathcal{X}_0, \mathcal{U} \subseteq \mathbb{R}^n$ are the initial states and the set of inputs given as compact convex sets.

5.2 Flowpipe Cost Computation

We define the cost of computing a flowpipe by considering a support function sample as the logical atomic task in the computation.

Definition 4. *Given a time horizon T, time discretization factor N, a set of template directions $\mathcal{D}$ and an initial symbolic state $s = (loc, \mathcal{C})$, the cost of computing the flowpipe with postC(s) is given by:*

$$flow_cost(s) = j.|\mathcal{D}| \begin{cases} j = max\{i \mid 0 \leq i \leq N, \forall 0 \leq k \leq i, \Omega_k \vdash Inv(loc)\} \\ \Omega_k \vdash Inv(loc) \text{ if and only if } \Omega_k \cap Inv(loc) \neq \emptyset \end{cases} \tag{5}$$

The longest sequence Ω_0 to Ω_j such that the convex sets satisfy the location invariant is identified. Since computing polyhedral approximation of the convex sets Ω requires sampling support function in each direction of the set of template directions $\mathcal{D}$, the $flow_cost$ essentially gives us the number of support function samplings, i.e. the atomic tasks, that is to be completed in order to compute the flowpipe. To compute $flow_cost$, it is necessary to find the longest sequence Ω_0 to Ω_j satisfying the location invariant $Inv(loc)$. Assuming polyhedral invariants, checking the invariant satisfaction can be performed using the following proposition.

Proposition 1. *[18] Given a polyhedra $\mathcal{I} = \bigwedge_{i=1}^{m} \ell_i \cdot x \leq b_i$ and a convex set Ω represented by its support function sup_Ω, $\Omega \vdash \mathcal{I}$ if and only if $-sup_\Omega(-\ell_i) \leq b_i$, for all $1 \leq i \leq m$.*

A procedure to identify the largest sequence is to apply Proposition 1 to each convex set starting from Ω_0 iteratively until we find a Ω_j such that $\Omega_j \nvdash Inv(loc)$. The time complexity of the procedure is $O(m \cdot N \cdot f)$, where f is the time for sampling the support function, m is the number of invariant constraints and N is the time discretization factor. We propose a cheaper algorithm with fewer support functions samplings for a class of linear dynamics $\dot{x} = Ax(t) + u$, with u being a fixed input. Fixed input leads to deterministic dynamics allowing to compute the reachable states symbolically at any time point .

Proposition 2. *Given an initial set $\mathcal{X}_0$ and dynamics $\dot{x} = Ax(t) + u$ with A being invertible, the set of states reachable at time t is given by:*

$$X(t) = e^{At}x_0 \oplus A^{-1}(e^{At} - I)(v) \tag{6}$$

The idea of the procedure shown in Algorithm 4, is to use a coarse time-step to compute reachable states using Proposition 2 and detect an approximate time for crossing the invariant. Once the invariant crossing time is detected, similar search is followed by narrowing the time-step for a finer search near the boundary of the invariant for a desired precision. The procedure is illustrated on a toy model of a counter clockwise circular rotation dynamics as shown in Figure 3a. The model has two locations with the same dynamics but different invariants. The transition assignment maps does not modify the variables. Figure 3b illustrates the procedure. The initial set on the location is shown in blue. The red sets are the reachable images of the initial set computed at coarse time steps to detect invariant crossing, followed by computing the images at finer time-steps shown in green near the invariant boundary for detecting an upper bound on the time of crossing the invariant with a desired precision. After computing this time, say t', the $flow_cost$ is obtained using Definition 4 with $j = t'/\tau$. However, the problem with the procedure is when it is possible for a reachable image to exit and re-enter the invariant within the chosen time-step. In such cases, the approximation error in the time returned by the procedure can be substantial. Constant dynamics and convex invariant $\mathcal{I}$ will not have such a scenario and the approximation error can be bounded.

Theorem 1. *For a class of dynamics $\dot{x} = k$, where k is a constant, let t be the exact time when reachable states from a given initial set $\mathcal{X}_0$ violate the convex location invariant $\mathcal{I}$. Let δ_C and δ_F be the coarse and fine time steps chosen to detect approximate time t' of invariant violation. The approximation error $|t - t'| \leq \delta_F$.*

Proof. Constant dynamics have a fixed direction of dynamics and therefore, convexity property ensures that the reachable states cannot exit and re-enter $\mathcal{I}$. Reachable states $X(t) = X_0 \oplus kt$ is exactly represented using its support function. Algorithm 4 samples the support function at δ_F time-steps to identify the time instant t' of crossing $\mathcal{I}$, which implies $|t - t'| \leq \delta_F$.

5.3 Discrete-Jump Cost Computation

The *PostD* computation performs the flowpipe intersection with the guard set followed by image computation. Considering a flowpipe having sets Ω_0 to Ω_j, each of these sets are applied with intersection operation and a map for non-empty intersection. Assuming intersection and image computation as the atomic task, the cost of *PostD* on a flowpipe $\cup_{i=0}^{j}\Omega_i$ will be j, which can be obtained from the $flow_cost$ computation in Definition 4. The addition of the cost of post operations for all symbolic states in the waiting list is used to uniformly

Algorithm 4 Detecting time of invariant crossing with varying time-step

procedure INVARIANT-CROSSING TIME DETECTION($\mathcal{I}$, $\mathcal{X}_0$, T)
  $discretization = 10$, $\tau = T/discretization$ ▷ Coarse Time-step
  $i = 0$, $\mathcal{R}(0) = \mathcal{X}_0$
  while $\mathcal{R}(\tau.i) \vdash \mathcal{I}$ **do** $i = i + 1$ ▷ Widened Search
  end while
  if $i > 1$ **then** $t1 = \tau * (i - 1)$
  else *return* 0
  end if
  $\tau = \tau/discretization$, $i = 0$ ▷ Fine Time-step
  while $\mathcal{R}(t1 + i * \tau) \vdash \mathcal{I}$ **do** $i = i + 1$ ▷ Narrowed Search
  end while
  return $t1 + i * \tau$ ▷ An upper bound on invariant crossing time
end procedure

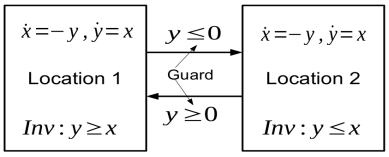

(a) Hybrid Automaton

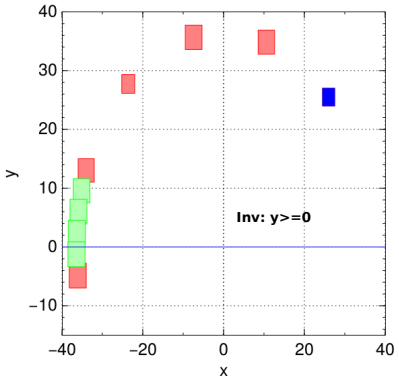

(b) Searching Time of Invariant Crossing with Widening-Narrowing Time-steps

Fig. 3. (a) A hybrid automaton of a toy example with circle dynamics. (b) Evaluation of the invariant crossing time detection algorithm on a circle model. The red sets are computed at coarse time-steps and the green sets computed at finer time-steps

distribute atomic tasks of post operations across the cores using multi-threading. Further details on the data-structures and task distribution is omitted due to the lack of space. The task parallel support-function-algorithm is referred as TP-BFS in the text that follows.

6 Experiments

The parallel algorithms are implemented with multi-threading using OpenMP compiler directives. Figure 5 shows the performance comparison between Reachability analysis with SpaceEx (LGG), Sequential BFS, A-GJH and TP-BFS. The

benchmarks are: A two dimensional oscillator circuit model, model of a bouncing ball under gravity, model of a circular rotation dynamics, three instances of the Navigation benchmark with 9 (3×3) locations and one instance each of 25 (5×5) and 81 (9×9) locations respectively. We conducted our experiments in a 4 core Intel i7-4770, 3.40GHz and 8GB RAM with hyper-threading enabled. The results are for a time horizon of 10 units, box template direction as parameters. The sampling time in Circle model is $1e-5$, in Oscillator, Timed Bouncing Ball, Navigation 3×3 and 5×5 instances are $1e-4$ and 0.1 units for Navigation 9×9 instance respectively.

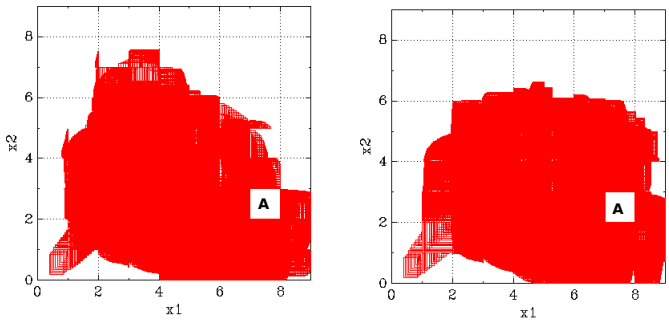

(a) XSpeed (b) SpaceEx(LGG)

Fig. 4. Reachable region of a 9×9 Navigation benchmark instance by (a) XSpeed after 13 BFS-levels (105563 post operations) and (b) SpaceEx's LGG algorithm (105563 post iterations)

In order to relate our results with SpaceEx (LGG) scenario, we apply the same number of post operations in both XSpeed and SpaceEx, with similar parameters. We count the number of post operations for a given bound on the BFS level in XSpeed and the exploration with SpaceEx is bounded with the same count on post operations (by setting the argument iter-max). However, note that the cost of post operations could be different for the symbolic states in XSpeed and in SpaceEx. Therefore, comparison on the number of post operations is not perfectly fair but we could not find a better means of comparing. Figure 4 shows that the reachable region obtained from XSpeed on a Navigation benchmark after 13 BFS levels (105563 posts) is comparable to that obtained from SpaceEx (LGG) after 105563 posts, and XSpeed computes the reachable region 900× faster, as shown in the Table.

The results for a Circle model is a good illustration of the effectiveness of the TP-BFS algorithm. BFS generates only two new symbolic states at every breadth, one of which exits the location invariants early leaving only one expensive flowpipe to be computed at each level. The A-GJH algorithm is slower than Seq-BFS algorithm due to the parallelization overhead. However, in case

Models	Breadths	Iter-max on SpaceEx	Time in Secs				CPU Utilization (%)			Speed-up		
			SpaceEx (LGG)	XSpeed (Seq-BFS)	XSpeed (A-GJH)	XSpeed (TP-BFS)	XSpeed (Seq-BFS)	XSpeed (A-GJH)	XSpeed (TP-BFS)	A-GJH vs. Seq-BFS	A-GJH vs. SpaceEx(LGG)	TP-BFS vs. SpaceEx(LGG)
Oscillator	2	2	4.78	0.67	0.79	0.53	12.5	12.7	31.0	0.8	6.1	9.0
	4	4	5.56	1.58	1.71	1.17	12.5	12.9	31.0	0.9	3.3	4.8
	6	6	6.48	2.14	2.56	1.73	12.5	13.0	31.2	0.8	2.5	3.7
Timed Bouncing Ball	2	2	12.65	2.91	3.34	1.43	12.5	12.6	58.7	0.9	3.8	8.8
	4	4	34.26	6.36	7.07	2.95	12.5	12.6	63.1	0.9	4.8	11.6
	6	6	196.01	9.22	10.11	4.22	12.5	12.6	63.0	0.9	19.4	46.5
Circle	4	7	37.68	29.39	33.92	13.66	12.5	12.5	61.9	0.9	1.1	2.8
	6	20	86.26	80.47	62.80	31.07	12.5	19.8	75.3	1.3	1.4	2.8
	8	54	217.48	213.08	133.47	77.12	12.5	33.8	82.1	1.6	1.6	2.8
Nav 3x3 (inst #1)	3	7	61.42	9.17	5.72	5.43	12.5	23.5	38.1	1.6	10.7	11.3
	5	35	199.27	44.37	24.24	18.87	12.5	39.3	65.5	1.8	8.2	10.6
	7	153	631.07	187.05	81.83	67.66	12.5	59.8	82.7	2.3	7.7	9.3
Nav 3x3 (inst #2)	3	6	181.19	15.53	12.13	8.83	12.5	18.4	36.8	1.3	14.9	20.5
	5	24	782.96	62.11	40.43	27.97	12.5	29.5	56.4	1.5	19.4	28.0
	7	91	4835.08	234.59	102.76	89.96	12.5	50.3	76.5	2.3	47.1	53.7
Nav 3x3 (inst #3)	3	7	133.88	10.94	7.98	5.79	12.5	19.5	43.7	1.4	16.8	23.1
	5	30	527.36	41.02	24.52	17.26	12.5	31.9	64.6	1.7	21.5	30.6
	7	109	1617.95	139.03	63.15	49.68	12.5	51.9	81.0	2.2	25.6	32.6
Nav 5x5	5	38	215.46	50.77	31.39	20.72	12.5	33.9	68.7	1.6	6.9	10.4
	7	194	1232.31	253.63	104.30	89.33	12.5	65.4	86.2	2.4	11.8	13.8
	9	980	6782.70	1290.86	449.96	429.59	12.5	81.0	90.1	2.9	15.1	15.8
Nav 9x9	5	46	0.70	0.09	0.05	0.06	11.1	50.0	72.9	1.6	13.3	11.9
	7	313	6.42	0.56	0.24	0.38	12.7	70.1	71.8	2.4	27.3	16.8
	9	2143	80.94	4.04	1.32	1.35	12.5	85.3	89.4	3.1	61.4	60.2
	11	14857	1687.06	28.77	8.43	10.50	12.5	91.9	85.3	3.4	200.2	160.7
	13	105563	56590.20	225.02	61.24	176.99	12.5	93.3	69.9	3.7	924.1	319.7

Fig. 5. Performance comparison of SpaceEx (LGG), sequential BFS, A-GJH and TP-BFS on hybrid systems benchmarks.

of TP-BFS algorithm the flowpipe tasks are distributed across all the available cores, making it faster than the Seq-BFS, A-GJH and SpaceEx (LGG).

We observe that when the number of explored symbolic states (shown as iter-max in Fig. 5) is low to moderate, TP-BFS shows better performance and CPU core utilization in comparison to A-GJH and SpaceEx (LGG). A maximum of 47.1× and 53.7× is observed on a 3×3 Navigation benchmark when 7 BFS levels are explored with a total of 91 symbolic states using A-GJH and TP-BFS respectively, with respect to SpaceEx's LGG scenario. We observe that when there is a large number of symbolic states in the waiting list, as in the NAV 9×9 instance, the CPU core utilization and performance of A-GJH is better than TP-BFS. We believe that this is because A-GJH keeps all available cores occupied, even if flowpipe computations are randomly assigned to cores, without taking their cost into consideration. In such a case, the extra overhead with task based load division becomes unnecessary as well as too expensive. This reduces CPU core utilization (since flowpipe cost computation and load-division is a sequential routine) and performance. This is illustrated in Fig. 6 which shows that the overhead of load balancing degrades the performance in TP-BFS with the increase in the explored symbolic states, whereas the A-GJH algorithm consistently gains in performance and utilization. We verified that for

the considered NAV 9×9 instance, the waiting list size in the BFS iterations are much larger than the available cores of the processor.

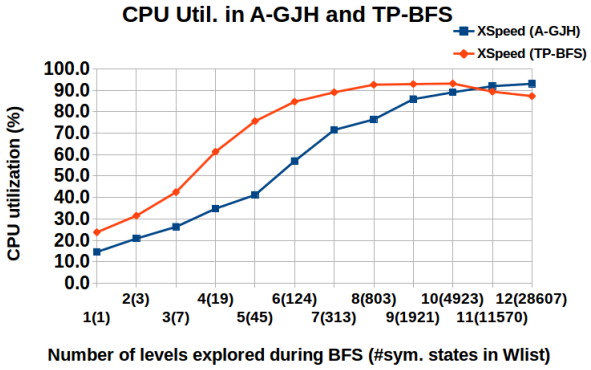

Fig. 6. Comparison of CPU Utilization between A-GJH and TP-BFS algorithm on a 9×9 Navigation model with sampling time as $1e-3$.

7 Conclusion

We present an adaption of G.J. Holzmann's breadth first exploration algorithm of the SPIN model checker, for reachability analysis of hybrid systems. We show that due to the intricacies of post operators in hybrid systems, this first approach does not always produce an efficient load balancing in the hybrid systems scenario. We then propose an alternative approach for load balancing that splits the tasks and distributes them evenly according to an efficiently precomputed cost of the post operations. We provide an implementation of this approach using a support-function based algorithm. Our experiments show that this second approach shows in general a better load-balancing and performance with respect to the first one, with the exception when the number of symbolic states to be explored in the next step is considerably very large. Overall, the two proposed algorithms show a considerable improvement in performance with respect to the current state of the art in reachability analysis for hybrid systems.

References

1. Bak, S., Bogomolov, S., Johnson, T.T.: HYST: a source transformation and translation tool for hybrid automaton models. In: Proc. of HSCC'15. pp. 128–133. ACM (2015)
2. Bartocci, E., Corradini, F., Berardini, M.R.D., Entcheva, E., Smolka, S.A., Grosu, R.: Modeling and simulation of cardiac tissue using hybrid I/O automata. Theor. Comput. Sci. 410(33-34), 3149–3165 (2009)

3. Bartocci, E., Lió, P.: Computational modeling, formal analysis, and tools for systems biology. PLoS Comput Biol 12(1), 1–22 (2016)
4. Bogomolov, S., Donzé, A., Frehse, G., Grosu, R., Johnson, T.T., Ladan, H., Podelski, A., Wehrle, M.: Guided search for hybrid systems based on coarse-grained space abstractions. International Journal on STTT pp. 1–19 (2015)
5. Bogomolov, S., Frehse, G., Greitschus, M., Grosu, R., Pasareanu, C.S., Podelski, A., Strump, T.: Assume-guarantee abstraction refinement meets hybrid systems. In: Proc. of HVC. pp. 116–131. LNCS, Springer (2014)
6. Bogomolov, S., Schilling, C., Bartocci, E., Batt, G., Kong, H., Grosu, R.: Abstraction-based parameter synthesis for multiaffine systems. In: Proc. of HVC. LNCS, vol. 9434, pp. 19–35 (2015)
7. Chen, X., Ábrahám, E., Sankaranarayanan, S.: Flow*: An analyzer for non-linear hybrid systems. In: Proc. of CAV 2013. LNCS, vol. 8044, pp. 258–263 (2013)
8. Chutinan, A., Krogh, B.H.: Verification of polyhedral-invariant hybrid automata using polygonal flow pipe approximations. In: Proc. of HSCC'99. LNCS, vol. 1569, pp. 76–90. Springer (1999)
9. Fehnker, A., Ivancic, F.: Benchmarks for hybrid systems verification. In: Proc. of HSCC. LNCS, vol. 2993, pp. 326–341. Sprnger (2004)
10. Frehse, G.: PHAVer: Algorithmic verification of hybrid systems past HyTech. STTT 10(3), 263–279 (2008)
11. Frehse, G., Le Guernic, C., Donzé, A., Cotton, S., Ray, R., Lebeltel, O., Ripado, R., Girard, A., Dang, T., Maler, O.: SpaceEx: Scalable verification of hybrid systems. In: Proc. of CAV. LNCS, vol. 6806, pp. 379–395. Springer (2011)
12. Girard, A.: Reachability of uncertain linear systems using zonotopes. In: Proc. of HSCC 2015. LNCS, vol. 3414, pp. 291–305. Springer (2005)
13. Girard, A., Le Guernic, C.: Efficient reachability analysis for linear systems using support functions. Proc. of IFAC World Congress 41(2), 8966–8971 (2008)
14. Guernic, C.L., Girard, A.: Reachability analysis of hybrid systems using support functions. In: Proc. of CAV 2009. LNCS, vol. 5643, pp. 540–554. Springer (2009)
15. Henzinger, T.A., Kopke, P.W., Puri, A., Varaiya, P.: What's decidable about hybrid automata? In: Journal of Computer and System Sciences. pp. 373–382. ACM Press (1995)
16. Holzmann, G.J.: Parallelizing the SPIN model checker. In: Proc. of SPIN 2012. LNCS, vol. 7385, pp. 155–171. Springer (2012)
17. Kong, S., Gao, S., Chen, W., Clarke, E.M.: dReach: δ-reachability analysis for hybrid systems. In: Proc. of TACAS 15. Lecture Notes in Computer Science, vol. 9035, pp. 200–205. Springer (2015)
18. Le Guernic, C.: Reachability analysis of hybrid systems with linear continuous dynamics. Ph.D. thesis, Université Grenoble 1 - Joseph Fourier (2009)
19. Lee, E.A., Seshia, S.A.: Introduction to Embedded Systems - A Cyber-Physical Systems Approach. E. A. Lee and S. A. Seshia, second edition edn. (2015)
20. Platzer, A., Quesel, J.: Keymaera: A hybrid theorem prover for hybrid systems (system description). In: Proc. of IJCAR. LNCS, vol. 5195, pp. 171–178. Springer (2008)
21. Ray, R., Gurung, A.: Poster: Parallel state space exploration of linear systems with inputs using xspeed. In: Proc. of HSCC'15. pp. 285–286. ACM (2015)
22. Ray, R., Gurung, A., Das, B., Bartocci, E., Bogomolov, S., Grosu, R.: XSpeed: Accelerating reachability analysis on multi-core processors. In: Proc. of HVC 2015. LNCS, vol. 9434, pp. 3–18 (2015)
23. Rockafellar, R.T., Wets, R.J.B.: Variational Analysis, vol. 317. Springer (1998)